

CLAIMS

What is claimed is:

1 1) In an Internet domain name server (DNS), a method of
2 providing an Internet Protocol (IP) address of a first one of a plurality of
3 servers of an Internet host, the method comprising the steps of:
4 receiving a DNS resolution request; and
5 returning the IP address of the first one of the plurality of servers
6 such that a total number of times that the IP address of the first one of the
7 plurality of servers is returned is proportional to a relative weight of the first
8 one of the plurality of servers relative to a total weight of the plurality
9 servers.

1 2) The method described in claim 1 including the additional
2 steps of:
3 periodically determining whether each one of the plurality of servers
4 is down;
5 setting to zero the relative weight of each one of the plurality of
6 servers that is down; and
7 recalculating the total weight of the plurality of servers.

1 3) The method described in claim 2 including the additional step
2 of informing the Internet host of each one of the plurality of servers that is
3 down.

1 4) The method described in claim 1 including the additional
2 steps of:
3 periodically determining whether each one of the plurality of servers
4 is generating a time out error;
5 reducing the relative weight of each one of the plurality of servers
6 that generate the time out error; and
7 recalculating the total weight of the plurality of servers.

1 5) The method described in claim 4 including the additional step
2 of notifying the Internet host of each one of the plurality of servers that is
3 generating the time out error.

1 6) The method described in claim 1 wherein the first one of the
2 plurality of servers is returned in a list including IP addresses of the other
3 ones of the plurality of servers such that a total number of times that the IP
4 address of the first one of the plurality of servers is at a top of the list is
5 proportional to the relative weight of the first one of the plurality of servers
6 relative to the total weight of the plurality servers.

1 7) In a domain name server (DNS), a method of providing an
2 Internet Protocol (IP) address of one of a plurality of servers of an Internet
3 host, the method comprising the steps of:
4 identifying a plurality of client domains that have recently requested
5 to be connected to the plurality of servers;
6 apportioning the plurality of client domains among the plurality of
7 servers such that a percentage of requests served by each one of the
8 plurality of servers is proportional to a relative weight of each respective
9 one of the plurality of servers;
10 receiving a DNS resolution request from one of the plurality of client
11 domains;
12 returning an IP address of one of the plurality of servers to which the
13 one of the client domains was apportioned.

1 8) The method described in claim 7 wherein the step of
2 apportioning the plurality of client domains among the plurality of servers
3 includes the step of identifying a number of times each one of the plurality
4 of client domains has recently requested being connected to the plurality of
5 servers.

1 9) The method described in claim 8 wherein the step of
2 identifying a number of times each one of the plurality of client domains has

3 recently requested being connected to the plurality of servers includes the
4 step of reading a recent server log of the Internet host.

1 10) The method described in claim 7 wherein the IP address of
2 the one of the plurality of servers to which the one of the client domains was
3 apportioned is returned in a list including IP addresses of the other ones of
4 the plurality of servers such that the IP address of the one of the plurality of
5 servers to which the one of the client domains was apportioned is at a top
6 of the list.

1 11) The method described in claim 7 including the additional
2 steps of:
3 periodically determining whether each one of the plurality of servers
4 is down;
5 setting to zero the relative weight of each one of the plurality of
6 servers that is down; and
7 recalculating the total weight of the plurality of servers.

1 12) The method described in claim 11 including the additional
2 step of notifying the Internet host of each one of the plurality of servers that
3 is down.

1 13) The method described in claim 7 including the additional
2 steps of:

3 periodically determining whether each one of the plurality of servers
4 is generating a time out error;
5 reducing the relative weight of each one of the plurality of servers
6 that generate the time out error; and
7 recalculating the total weight of the plurality of servers.

1 14) The method described in claim 13 including the additional
2 step of notifying the Internet host of each one of the plurality of servers that
3 is generating the time out error.

1 15) The method described in claim 7 including the additional
2 steps of:
3 receiving a DNS resolution request from a client domain not
4 included in the plurality of client domains that have recently requested to be
5 connected to the plurality of servers; and
6 returning the IP address of a first one of the plurality of servers such
7 that a total number of times that the IP address of the first one of the plurality
8 of servers is returned is proportional to a leftover capacity of the first one of
9 the plurality of servers relative to a total leftover capacity of the plurality
10 servers.

1 16) The method described in claim 7 wherein the plurality of client
2 domains that have recently requested to be connected to the plurality of
3 servers include client domains that have recently frequently requested to
4 be connected to the plurality of servers.

1 17) In a domain name server (DNS), a method of providing an
2 Internet Protocol (IP) address of one of a plurality of servers of an Internet
3 host, the method comprising the steps of:
4 determining a relative weight for each one of the plurality of servers;
5 calculating a total weight of the plurality of servers
6 identifying a plurality of client domains that have recently frequently
7 requested to be connected to the plurality of servers;
8 identifying a number of times each one of the plurality of client
9 domains has recently connected to the plurality of servers
10 calculating a total number of times the plurality of servers has been
11 connected to the plurality of client domains;
12 apportioning the plurality of client domains among the plurality of
13 servers such that the relative weight of the first one of the plurality of servers
14 in comparison to the total weight of the plurality servers is approximately
15 proportional to a cumulative number of times each one of the plurality of
16 client domains apportioned to the first one of the plurality of servers has
17 recently connected to the plurality of servers in comparison to the total

18 number of times the plurality of servers has been connected to the plurality
19 of client domains;
20 receiving a DNS resolution request from one of the plurality of client
21 domains; and
22 returning an IP address of one of the plurality of servers to which the
23 one of the client domains was apportioned.

1 18) The method described in claim 17 including the additional
2 steps of:
3 receiving a DNS resolution request from a client domain not
4 included in the plurality of client domains that have recently frequently
5 requested to be connected to the plurality of servers; and
6 returning the IP address of a first one of the plurality of servers such
7 that a total number of times that the IP address of the first one of the plurality
8 of servers is returned is proportional to a leftover capacity of the first one of
9 the plurality of servers relative to a total leftover capacity of the plurality
10 servers.

1 19) The method described in claim 17 including the additional
2 steps of:
3 periodically determining whether each one of the plurality of servers
4 is down;

5 setting to zero the relative weight of each one of the plurality of
6 servers that is down; and
7 recalculating the total weight of the plurality of servers.

1 20) The method described in claim 17 including the additional
2 steps of:
3 periodically determining whether each one of the plurality of servers
4 is generating a time out error;
5 reducing the relative weight of each one of the plurality of servers
6 that generate the time out error; and
7 recalculating the total weight of the plurality of servers.